UTAH DEPARTMENT OF TECHNOLOGY SERVICES

August 2010

Data Center Consolidation Initiative

Table of Contents

Overview: Data Center Consolidation Initiative

The State of Utah Department of Technology Services (DTS) has implemented a Data Center and Server Consolidation. The initiative included the consolidation of data centers and servers throughout the Executive Branch agencies in the State. The project was successfully completed in a short 12 months, with a 6 month planning period beforehand. The results of the project include:

- Reduced the number of data centers in the State from 35 to 2
- Reduced the total number of physical servers from 1864 to 591
- Save the State $4 million annually ongoing in server, support, and energy costs
- Increased security
- Increased performance

Performance throughout State applications has improved significantly with the completion of the consolidation. The consolidation has allowed resources to be broken into smaller, more flexibly used components and pooled for more effective utilization. Below are a few of the many instances where consolidation and virtualization has increased performance, allowing the State to become more efficient, improve business practices, and provide better service to citizens:

- 60% performance gain on a Data Warehouse batch job that is regularly run at the Department of Workforce Services
- Perceivable 30% performance increase at the Department of Health when accessing data files with outside entities
- Total runtime for State Payroll decreased from 39 hours to just 3.5 hours, and realized a cost avoidance of $300,000 in hardware

Project Summary

The Data Center Consolidation project for the State of Utah commenced on January 6th, 2009. It was imperative to initially gain the support of the Executive Branch Agencies, the Governor, and the Legislature in order to get the buy-in of the business and to also increase the accountability of DTS to report regularly to sponsoring entities. Three Project managers were chosen to oversee the project, each with differing IT backgrounds and experiences. In the initial kickoff meeting, the following four objectives were presented and discussed:

- Reduce the data center count within the Utah Executive Branch from 35 to 2
- Reduce the number of servers in the 35 data centers from 1700 to 400
- Achieve an ongoing yearly savings of 4 million dollars.
- Complete the project by June 30, 2010

Planning involved meeting with vendors and states who either had conducted their own data center consolidation initiatives internal to their company or state, or had consulting expertise in successfully planning and implementing a data center consolidation project of our size and scope, or who sold hardware/software that could be used in such a consolidation effort. Specifically, IBM, HP and Sun Microsystems were a few of the many companies who worked with DTS, offering either free consulting or consulting at extremely reduced rates to help in the planning process. Both the states of California and Oregon also offered valuable feedback on previous and current consolidation efforts.

To help manage the project, it was decided to manage at the level closest to the customer. All of the 35 data centers existed in one of seven campuses, with hosting personnel assigned to work in each of those campuses. A project lead was identified in each of those seven campuses with specific instructions to develop a plan to virtualize and move systems in their campus, to one of the two enterprise data centers located in Salt Lake City or Richfield.

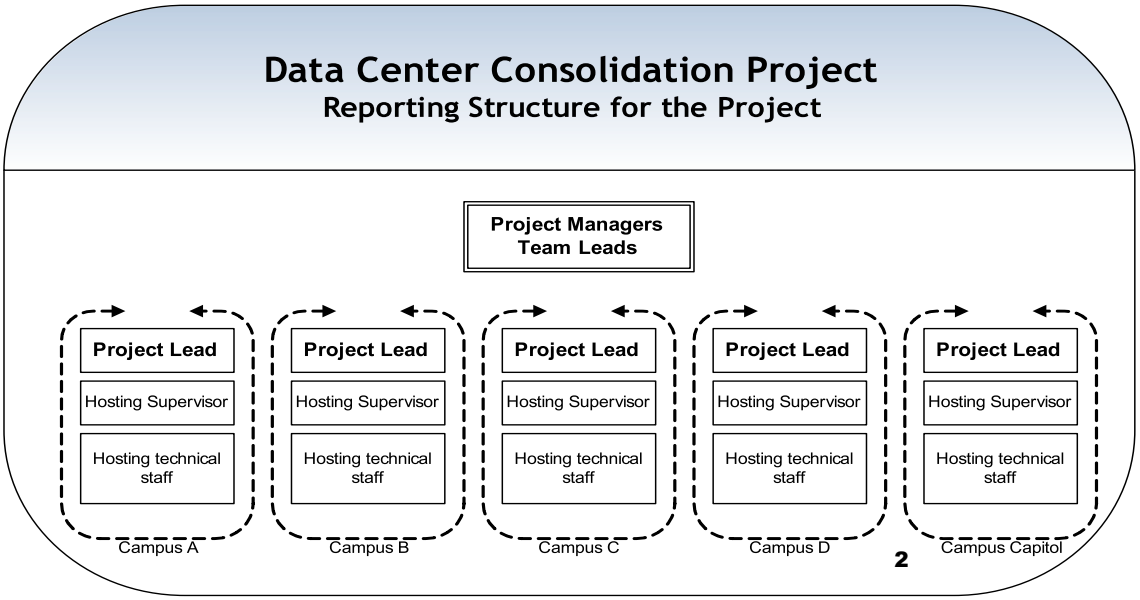

By having a project lead at each campus, the project was managed at the level closest to the business and gave the seven project leads the ability to focus on issues unique to the customers in their respective campuses. It also allowed the Project Managers overseeing the entire project to manage the project more like a program made up of multiple projects, rather than one large project. ***This allowed the campuses to make progress independent of the other campuses and allowed the campuses to work concurrently rather than consecutively, which dramatically reduced the amount of time needed to complete the project.***

Along with having project teams at each campus focused on doing consolidation and virtualization of servers and server migration, project teams were established surrounding enterprise operational needs of hosting, networking, storage and backups, and virtualization. Each of these teams had one representative from each campus on the team so that the campus would have their operational concerns addressed as the technical architecture was planned for and implemented.

March 2009 through December 2009, decisions were made on hardware, software, and consulting services. The following was decided:

Hosting: A decision was made to standardize the virtual platform on HP hardware using HP Blade Centers (c7000) and blades (BL460c G6). Each chassis was populated with 16 blades. Each blade was populated with 96 GB of RAM.

Virtualization: The decision was made to virtualize using VMware. Because the decision was made while VMware was in the midst of updating its product line from 3.x to VSphere 4.0, and because of concerns with deploying a dot 0 product, DTS elected to stay with the 3.x product for this project, with the intent of migrating to 4.x after the project was complete.

Backups: DTS already had a significant deployment of IBM's Tivoli Storage Manager in their previous environment, along with experience in using that product. Because of this, the team decided to continue using the Tivoli product. Continuing to use Tivoli Storage Manager, which was already a sunk cost, saved the expense of purchasing a replacement backup solution and lessened the time needed to deploy a similar competitor's solution. To help lower costs of storing backed up data, DTS deployed a de-duplication storage product from Data Domain (dual DD880 controllers, 140 TB of storage in each controller, one controller deployed in SLC and Richfield Data Centers). The Data Domain solution was implemented as virtual tape drives and wrote the data directly to disk, speeding up both the backup and the restore processes. As an archival option, DTS can choose to move backups to tape based on the number of days residing on disk.

Storage: The team decided to deploy a Sun 9990v (rebranded Hitachi USP-VM) in each of the data centers with modular storage (Sun Storage 6780s and Hitachi AMS 2500) behind each of the controllers. This allowed the team to virtualize the storage and save space and money. Now, instead of having allocated storage taken out of the allocation pool when it is initially requested by the system administrator, storage space is marked for use but is not completely committed to the system until the system actually uses the storage. This allows for storage architects to intentionally over subscribe storage space. By tracking actual use of the storage, the storage architect can still plan for expansion without being forced to expand because storage was requested but not used. The SAN fabric is dual Cisco MDS 9513 with 4/8 GB capability. The management of the system is conducted using Tuning Manager and Tiered Storage Manager. Tiered Storage Manager gives us the ability to migrate data from different tiers without any disruption to the end user.

Networking: The hosting area was expanded to handle the influx of servers to the Salt Lake Data center. Two new hosting rows were built, each with a Cisco Nexus 5010 24 port 10 Gig switch. Each of the HP Blade Chassis was dual homed to each of these switches. ACE modules that already existed in the core network infrastructure were provided to systems that required the network, so that failover could occur for systems that needed redundancy. Firewall Service Modules were also scaled so that each agency could maintain their own firewall instance, thereby allowing the agency to continue being the custodian of their own data without having to negotiate with other agencies over access privileges or security requirements of the data. Additionally, a second redundant data path from the state WAN was installed in the Salt Lake Data Center (SLDC). The second line was given a separate physical path than the first line. Finally, a second Internet connection from a different service provider was installed to offer further redundancy to the SLDC with all internet traffic being load balanced across both internet connections.

As DTS worked through the planning process with each of the campuses, which included verifying the number of servers and data centers, it was discovered that there were actually 1864 servers and not 1700, an increase of 9% over the initial server count. Fortunately, the increased server count was discovered early enough that DTS was able to scale the architecture to account for the increase.

The actual work of physically migrating systems from the 35 data centers to the Salt Lake and Richfield Data Centers started taking place in earnest in January 2010. Over the life of the project, the following is what was reported by each campus:

Utah Data Center Consolidation Project Jan 2009 - June 2010								
	Campus A	Campus B	Campus C	Campus D	Campus Capitol	Campus CR	Campus Rural	All Total
Total Servers (Original)	316	239	248	233	736	54	38	1864
Data Center Completed	4 of 4	8 of 8	5 of 5	5 of 5	5 of 5	2 of 2	6 of 6	35 of 35
Servers Completed	316	239	248	233	736	54	38	1864
Physical Servers (Final count)	71	72	78	43	224	25	9	591

Below is a breakout of the Data Center Consolidation Project by week. The numbers on the left reflect the number of servers that were dealt with in the project. The numbers on the bottom reflect the weeks in the project. Week one occurred the first week of November 2009. Week 37 occurred the last week in June 2010. The straight blue line is the trend line that needed to be met if the project was to be finished on time. The straight red line is the trend line of the actual consolidation work done for the project. The thick red line represents the actual work done for each week of the project.

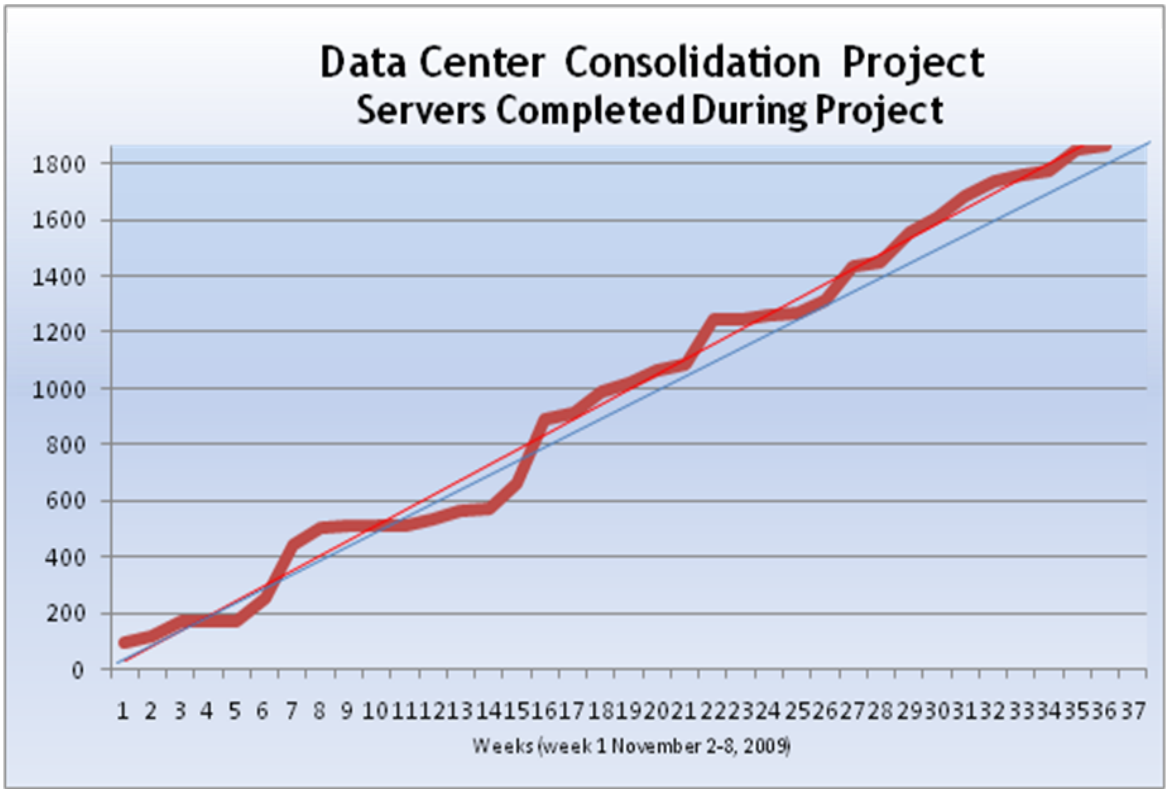

Accomplishments

The Data Center Consolidation Project was able to realize the following overall accomplishments:

1. Successfully consolidated 35 existing data centers down to 2 data centers.

2. Successfully reduced physical server counts from 1864 servers, to 591 physical servers, a 68% reduction in server counts. In two different consulting engagements, the State of Utah was told that as an industry average, 20-30% of the servers in a consolidation effort are typically not able to be virtualized and moved, virtualized only, or moved only. The project came close to meeting the industry average even though a decision was made to not move database servers into a shared virtual environment because of concerns over costs of licensing databases in that environment.

3. Built out a scalable enterprise storage environment that allows for Tier 1, Tier 2, and Tier 3 data storage. The storage is flexible enough to allow for modular storage from multiple vendors, yet robust enough to handle even the most demanding applications.

4. Successfully trained the majority of system administrators within Utah Department of Technology Services (20-25 employees) to effectively use the Hosting Environment's virtualization hardware and software. These employees had very little knowledge of this technology or used it minimally before the project began.

5. The network team for the state successfully configured and implemented MPLS on the state's core network. To ensure that communication to systems relocated from the agency sites would not be inhibited by non-critical internet traffic, QOS was configured at many of the agency sites that used to have the servers housed locally.

At a Campus level, the Data Center Consolidation Project was able to realize the following accomplishments:

Campus A (DWS, Commerce, Labor Commission, DWS e-REP American Fork)

Measure	DWS - Main	Com-merce	Labor Commis sion	DWS - eREP Americn Fork
Delivery Target	6/22/10	11/20/09	10/24/09	9/21/09
Total Servers Completed	276	10	10	20

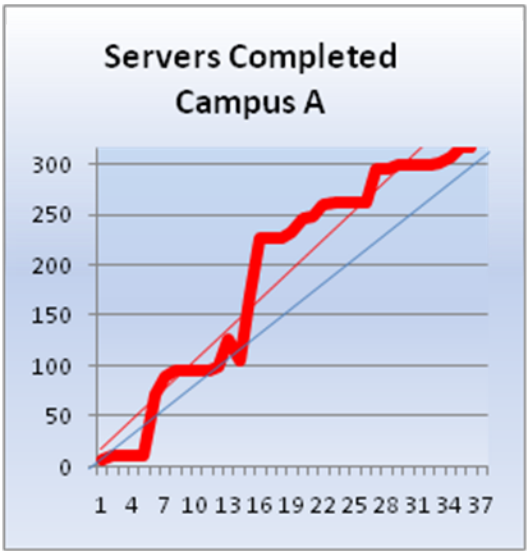

In preparation for the server moves, the hosting staff assigned to Campus A successfully segmented a network at the main Department of Workforce Services building consisting of 300+ servers into four different networks without impacting the agency. While this was a necessary configuration change to ensure a successful migration of DWS servers, the reconfiguration had the added benefit of alleviating some recurring random network problems on the local area network.

With the help of Network Operations, the technical staff successfully implemented MPLS and Psuedowire so that a DWS 50 server migration for DWS could be completed over a over a four week time period instead of one night. By implementing Psuedowire, the subnet for DWS CUBS, CATS, and E-work existed at both DWS and the SLDC so that as servers moved, they continued communicating with servers as if they were all located at the same physical location. This allowed for a more methodical move of these mission critical systems and met the requirements of the business.

When DWS started out with the Data Center Consolidation project, there were 278 servers. After many of those server instances were migrated to the shared hosted environment being used by all of the other agencies, DWS was left with 71 physical servers. 30 of the Sun Microsystems systems were consolidated down to 15 Sun Boxes or Linux instances in the shared environment, a significantly less costly alternative to Sun. In some instances, there were significant performance gains in the migration. One system administrator reported a 60% performance gain on a Data Warehouse batch job that is regularly run.

Much of the DWS data center migration occurred during quarter end for Unemployment. Even though it was a record-breaking quarter for the submittal of new and recurring claims for the state, the migration was able to occur with very little impact to the agency.

Campus B (DHS, Tax, DEQ, DCC)

Measure	State Library	324 S. State	DEQ	DHS Admin	ORS	Tax	USDC	USH
Delivery Target	2/6/10	11/30/09	4/30/10	11/30/09	2/28/10	6/30/10	5/30/10	6/30/10
Total Servers Completed	6	1	37	64	14	99	7	11

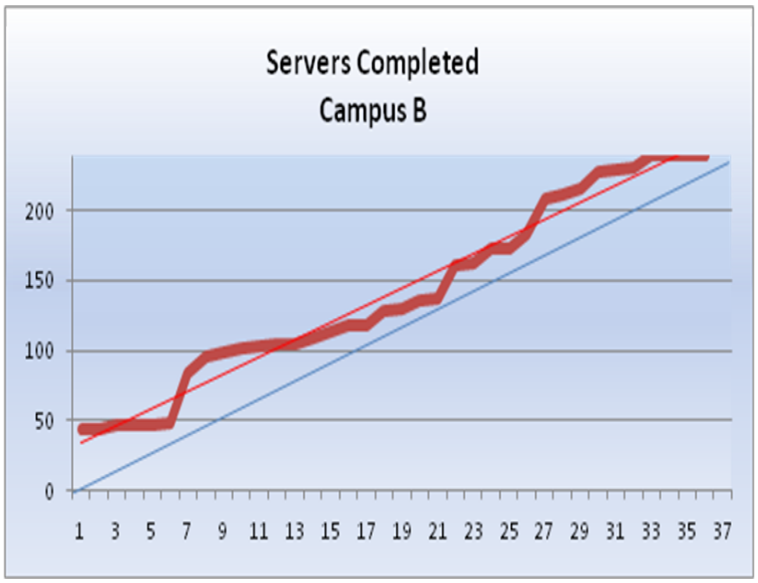

The major applications for the Campus B data centers were successfully moved to the SLDC with little or no downtime experienced by the majority of customers. Some of the notable systems moved include:

- SAFE (DHS)
- Motor Vehicles (Tax)
- Gentax (Tax)
- Documentum (DEQ)
- KLAS
- Webgrants
- eChart

Approximately 167 servers from Campus B were eliminated by moving system to the shared environment and consolidating applications onto the same servers.

Campus C (UDOT, UDOT TOC, UDOT Region 2, DPS Rampton, DPS POST)

Measure	UDOT Calvin Rampton Center	UDOT TOC	UDOT Region 2	DPS Calvin Rampton Center	DPS Post
Delivery Target	3/31/10	5/15/10	5/15/10	5/15/10	5/15/10
Total Servers Completed	153	47	8	36	4

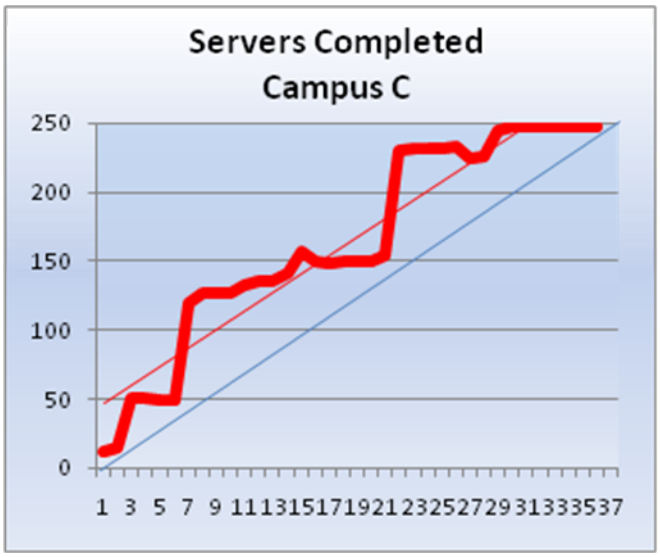

Campus C successfully planned and coordinated a migration of 62 Virtual servers and 12 physical servers to the Salt Lake Data Center on the same day. Falconstor replication technology was used to facilitate the migration of virtual machines and data to the Salt Lake Data Center prior to the cut over weekend. Planning meetings included Campus Infrastructure staff, Application Development staff and key business customers. Each of these separate groups had the responsibility to perform testing to ensure that the applications worked as expected before and after they were moved. The move started on Saturday March 20th, at 6 AM and by 9:30 PM, the application development group had certified the Application Developers as well as the business customers. Because of the excellent planning and implementation of the Campus, the migration was virtually transparent to the customer.

Campus D (Health, DNR, ABC, DOH 44 Med, Agriculture)

Measure	DOH and CHB	DABC	DNR	DOH and 44MED	Agriculture
Delivery Target	5/30/10	5/23/10	5/30/10	5/30/10	12/31/09
Total Servers Completed	143	7	68	9	6

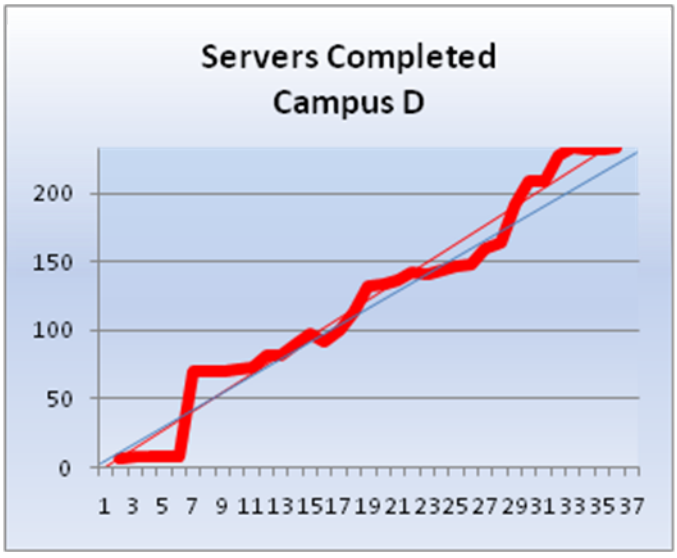

Campus D began the project with a target of over 250 servers spread across 4 Datacenters. 96 servers were retired or their services were consolidated to other servers, 130 servers were virtualized then moved to the SLDC, 21 servers that could not be virtualized were physically moved to the SLDC and 20 new servers were created at the SLDC which resulted in a significant hardware savings for the Agency Programs.

Campus D successfully moved three "N-Tiered"systems that supply remote applications to users throughout the state via P2P-VPN connections with outside entities for bi-directional message exchange. These services were successfully re-created at the Salt Lake Data Center, and provided a nearly seamless transition from the VPN concentrator at the Canon Health Building. Reports received from the users affected by these moves indicated a perceivable 20-30% performance increase accessing data files, and an overall application performance improvement.

Campus Rural (Ogden Regional Center, Provo Regional Center, Clearfield West, UDOT Regions 1,3,4)

Measure	Ogden Regional Center	Provo Regional Center	Clearfield West	UDOT Region 1	UDOT Region 3	UDOT Region 4
Delivery Target	5/30/10	2/28/10	11/30/09	5/1/10	5/1/10	5/1/10
Total Servers Completed	8	6	4	6	6	8

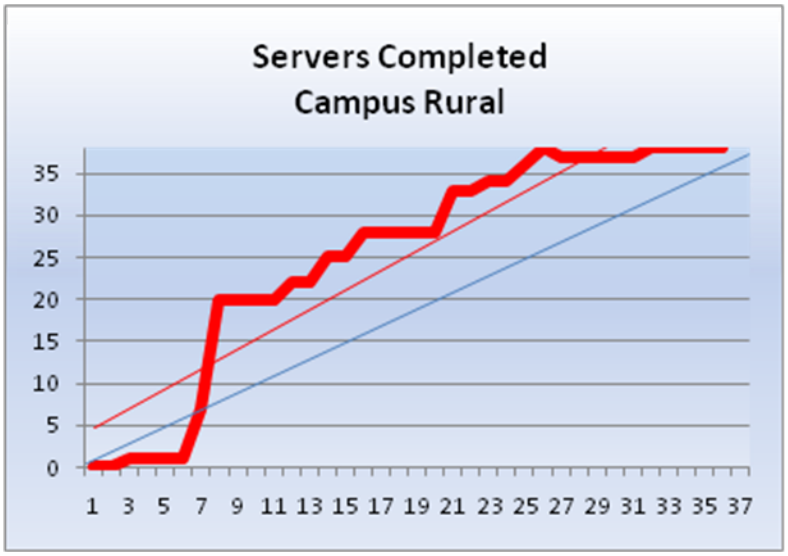

The rural campus discovered 8 servers in the rural campus that were being used as internet filters. By enabling reporting for the enterprise internet filtering already in place, the customers agreed to turn off the internet filter they were paying for. This will result in a cost savings for the state as redundant services like this that are no longer needed are eliminated.

UDOT region offices in Orem and Ogden had six servers and five servers respectively. When consolidation and virtualization was completed, each location had one server running one virtual host with four or five server instances (Windows and Novell), and one backup server.

Campus Capitol (all servers in SLDC and on Capitol Hill)

Measure	Salt Lake Data Center	Governor's Office	DHRM	Division of Finance	Richfield Data Center
Delivery Target	6/30/10	6/30/10	8/30/09	4/30/10	6/30/10
Total Servers Completed	609	17	23	14	73

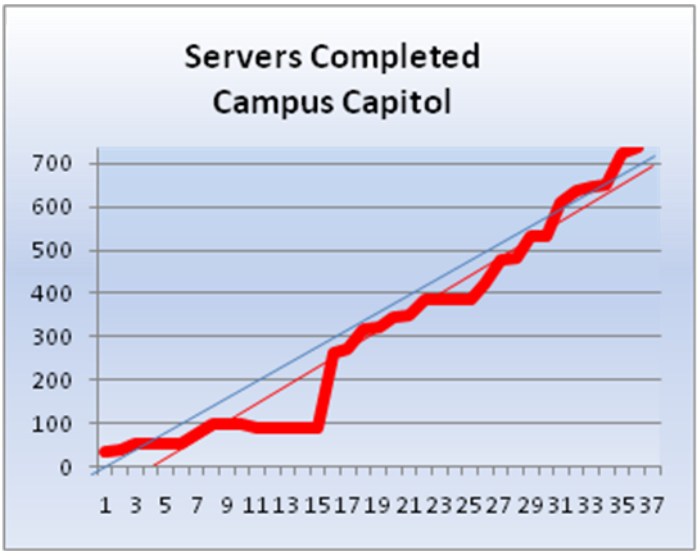

Working with the Division of Finance, SAP was migrated from an HP-UX hardware platform to a SuSE Linux VM environment. This saved the Division of Finance from purchasing new PA-RISC or Itanium HP hardware and has shown improvement in the performance of batch processes as well as the Employee Self-Service (ESS) portal for time-entry. Total runtimes decreased from over 39 hours to just 3.5 hours.

Campus Corrections (Point of the Mountain, Gunnison)

Measure	Draper Data Center	Gunnison
Delivery Target	3/31/10	3/31/10
Total Servers Completed	50	4

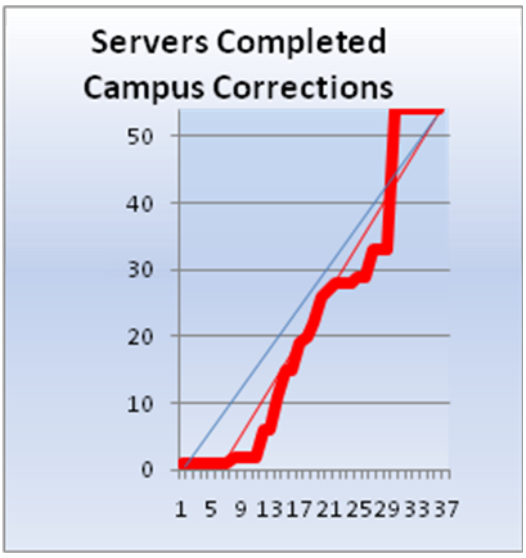

The Corrections technical staff had moved to a virtual computing environment on HP blade servers a few years ago meaning they were already versed in the technology required for the consolidation effort. Nearly half of the 53 servers were already in a virtual environment giving the staff an advantage in the task of moving the servers to the SLDC virtual infrastructure. The biggest challenge was in coordinating individual server moves with Correction business units, units that many times had 24/7 uptime requirements. Coordination of server moves was, by far, more challenging than the technical aspect of the migrations.

Actual expenses - final budget

DATA CENTER CONSOLIDATION PROJECT							
DTS One-Time Costs							
Data Center Project Purchases		Anticipated Cost	Budget $ Remaining	Budget $ Spent	DC	Purchases	Actual Cost
(DC1)Hosting (X86)							
	Proposed purchases				1.1	HP/IBM Blade Systems	$1,697,664
DC1.1	7 Blade Centers @ $230,000 (112 Blades)	$1,091,063			1.2	3 DL380 Metro assist	$62,248
DC1.1a	Additional Blade Centers	$465,040			1.4	SAS drives	$3,536
DC1.2	3 DL 380 servers to assist Agency's in Virtualization	$37,276			1.5	RAM	$20,315
DC1.3	3 DL 380 servers to assist Rural's Virtualization	$32,000			1.6	Acronis	$1,614
DC1.4	Extra SAS hard drives to aid in virtualization at	$6,000			1.3	3 Rural Virtualization Servers	$33,032
DC1.5	RAM for Data Warehouse server (DWS)	$15,645					
DC1.6	Disk Imaging software Acronis (Capital)	$1,614					
DC1.1b	Additional Blade Centers	$132,500					
	Total	$1,781,138	-$208,409	$1,818,409			
(DC2)Monitoring							
	Total budget				2.1	Tuning Manager	$13,044
DC2.1	Tuning Manager software for SAN Storage	$13,943				up.Time CP Software	$101,523
DC2.2	Capacity Planning and Monitoring Software	$100,000					
DC2.3							
	Total	$113,943	$45,433	$114,567			
(DC3)(Sun,AIX, etc) /Agency Hardware							
	Total budget				3.3	Equipment trade in compensation to Agencies	
DC3.1	SUN	$300,000			3.3	HP trade-in credit	-$19,018
DC3.2	AIX	$230,000			IN4.6	Phase III Storage (Excess budget spent from this category to aquire additional storage)	$1,095,029
DC3.3	Agency Hardware buy back	$300,000					
	Total	$830,000	-$246,011	$1,076,011			
(DC6)Virtualization (Vmware)							
	Total budget				6.1	3 year contract for Vmware	$897,491
DC6.1	VMWare Enterprise License	$900,000			6.2	VmWare Heartbeat	$14,926
DC6.2	VMWare vCenter Server Heartbeat	$11,081					
	Total	$911,081	-$12,417	$912,417			
(DC7)Developer/Consultant Costs							
	Total budget				7.1	Finance SAP	$14,000
DC7.1	SAP Finance	$14,000			7.2	TSM Consultant (TEK Systems)	$139,400
DC7.2	Sys Admin to augment Agency Staff (Tek	$340,900			7.2	VM Consultant (Compunet)	$173,700
DC7.3	Agency Development Cost	$275,000			7.3	Vendor Hardware moves assist.	$22,106
DC7.5	Vendor Hardware moves assist.	$4,089			7.6	Core Team Travel	$3,190
DC7.6	Core Team Travel	$3,190					
DC7.2	Consultant Costs	$38,200					
DC7.3	AMS 500 Maintenance (TAX)	$10,800					
	Total	$686,179	$147,604	$352,396			
	Total	$4,322,341	-$273,800	$4,273,800			

Successful Practices and Lessons Learned

- Sponsorship of Governor and Legislature is essential to the success of the project and the buy-in of the business
- Support of executive management is essential to the success of the project and provides a means to be accountable for reporting
- Ensure the project plan has clearly defined objectives and deliverables to prevent scope creep
- Give visibility to the progress of the project
- Break up the project into manageable parts and run the project as concurrent multiple projects instead of one consecutive project
- Give each project group the responsibility to form their own plan and accountability for achieving the architected plan
- Create a CMDB of all equipment impacted by the project the first step in the project
- Regular communication is essential for a successful project
- Regular and consistent communication needs to exist between enterprise groups (e.g hosting, storage, networking) and IT staff assigned to the campus
- Ensure teams are being heard
- Customer Communication - Maintain a high level of communication with the customer and the Campus IT manager about the project
- Ensure there are single points of contact for each of the functional areas of networking, storage, virtual environment, and data center
- Don't rely on email as the sole source for communicating changes and requests for the project
- Ensure funding levels are appropriately identified for the project
- Ensure technical groups are committed to the project
- Ensure the organizational structure is aligned with project
- Identify processes of moving the hardware and communicate those moves to the enterprise groups

- Understand that even with good sound processes and adequate planning, tasks are going to be forgotten in a large-scale move, and technical teams will need to possess the agility to react to the situation

- Incident, problem, change and reliability management tools and processes should be in place prior to project initialization to prevent service level issues

- Not every campus can or should do their moves like the other campuses

- Ensure the enterprise group has the hardware and people in place to handle the scaling of systems and resources that will need to occur

- Virtualization is not a one solution fits all

- With complex systems, it was easier to P2V than to migrate them. With simplistic systems, it was easier to migrate than to P2V.

- Set time frames for task completions to allow for application testing before/after server migrations

- Ensure that staff have knowledge of a virtual server environment

- Judicious use of consultants to fill in the knowledge gaps

- Recognize the hard work of staff